# Volunteer Computing: the Ultimate Cloud

Dr. David P. Anderson
University of California, Berkeley

Oct 19, 2010


# Is Amazon EC2 a good deal?

- Yes, for sporadic or unpredictable workloads
  - e.g., a Super Bowl web site
- No, for a constant processing workload
  - e.g., many types of scientific computing

# The yearly cost of 10 TeraFLOPS

- Amazon EC2
  - small instance: $.09/hour = $788/year
  - 10 TeraFLOPS = 5,000 instances
  - $3.94M/year plus network, storage costs
- Build your own cluster
  - ~ $1.5M/year
- Volunteer computing
  - ~ $0.1M/year

# Volunteer computing

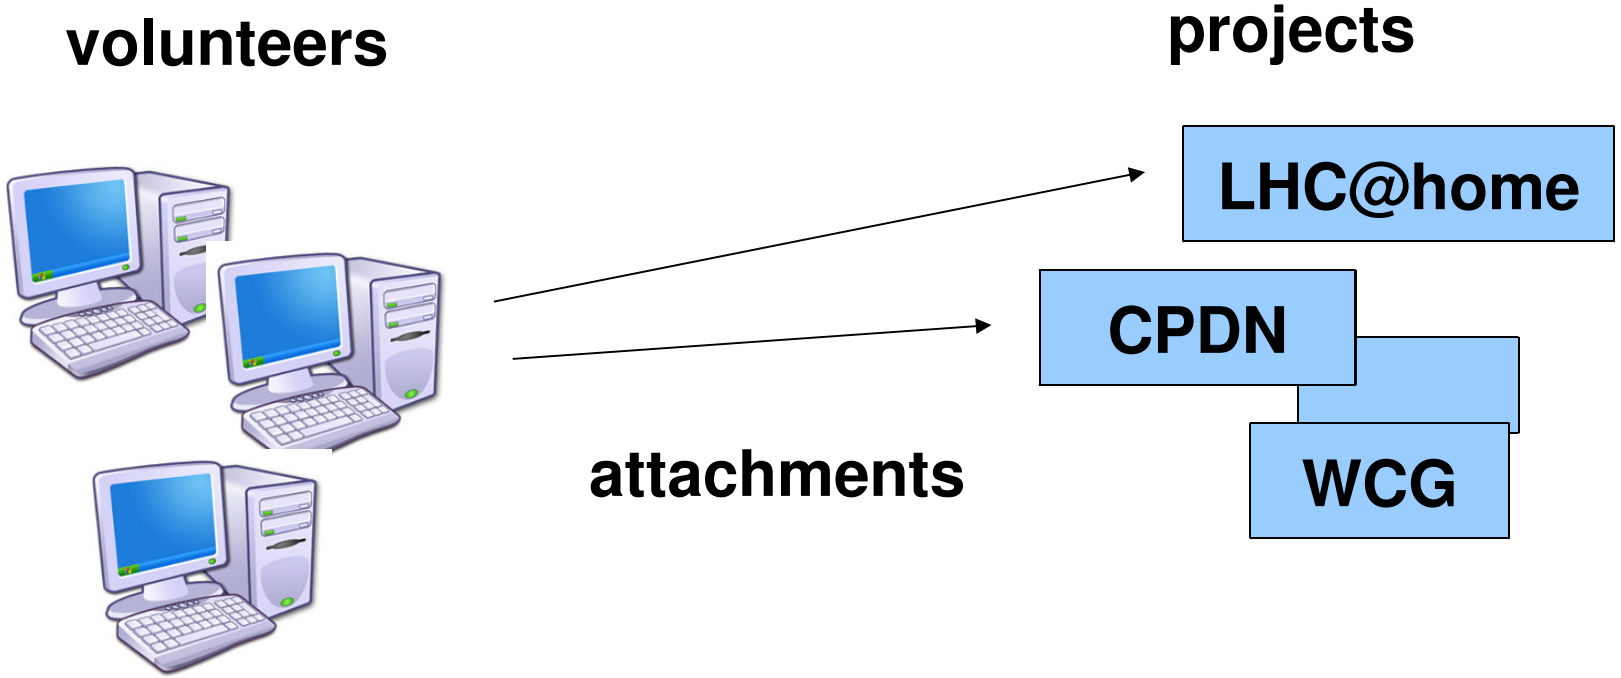


- Scientists create projects using BOINC
- Volunteers install BOINC, attach to project(s)
- Applications are silently downloaded and executed on volunteer PCs

# The Consumer Digital Infrastructure

- 1.5 billion PCs
- Graphics Processing Units: TeraFLOPS
- Terabyte-scale storage
- Network speed approaching 1 Gbps
- Ideal for scientific computing!

# Consumer versus Institutional computing resources

- Capacity
  - Institutional:  ExaFLOPS supercomputer in 5 years?
  - Consumer: ~1000 ExaFLOPS today
- Cost
  - Institutional: ~$200M/year from funding agencies
  - Consumer: ~$1 trillion/year from public, self-replenishing, self-maintaining, self-powering

# The state of volunteer computing

- 40 projects
- 500K volunteers
- 800K computers
- 10 PetaFLOPS
  - would cost $3.94 billion/year on Amazon EC2

# Science areas using BOINC

- Biology: protein study, genetic analysis
- Medicine: drug discovery, epidemiology
- Physics: LHC, nanotechnology, quantum computing
- Astronomy: LIGO, radio data analysis; cosmology; galactic modeling
- Environment: climate modeling, botanical ecosystem simulation
- Math

# Organizational issues

- Single-scientist projects: a dead end
  - Barriers to entry are too high
  - Wrong marketing model
  - Doesn’t handle sporadic requirements
- Umbrella projects
  - IBM World Community Grid
  - Campus-level (UCBerkeley@home)

# A better model: ScienceUSA.org

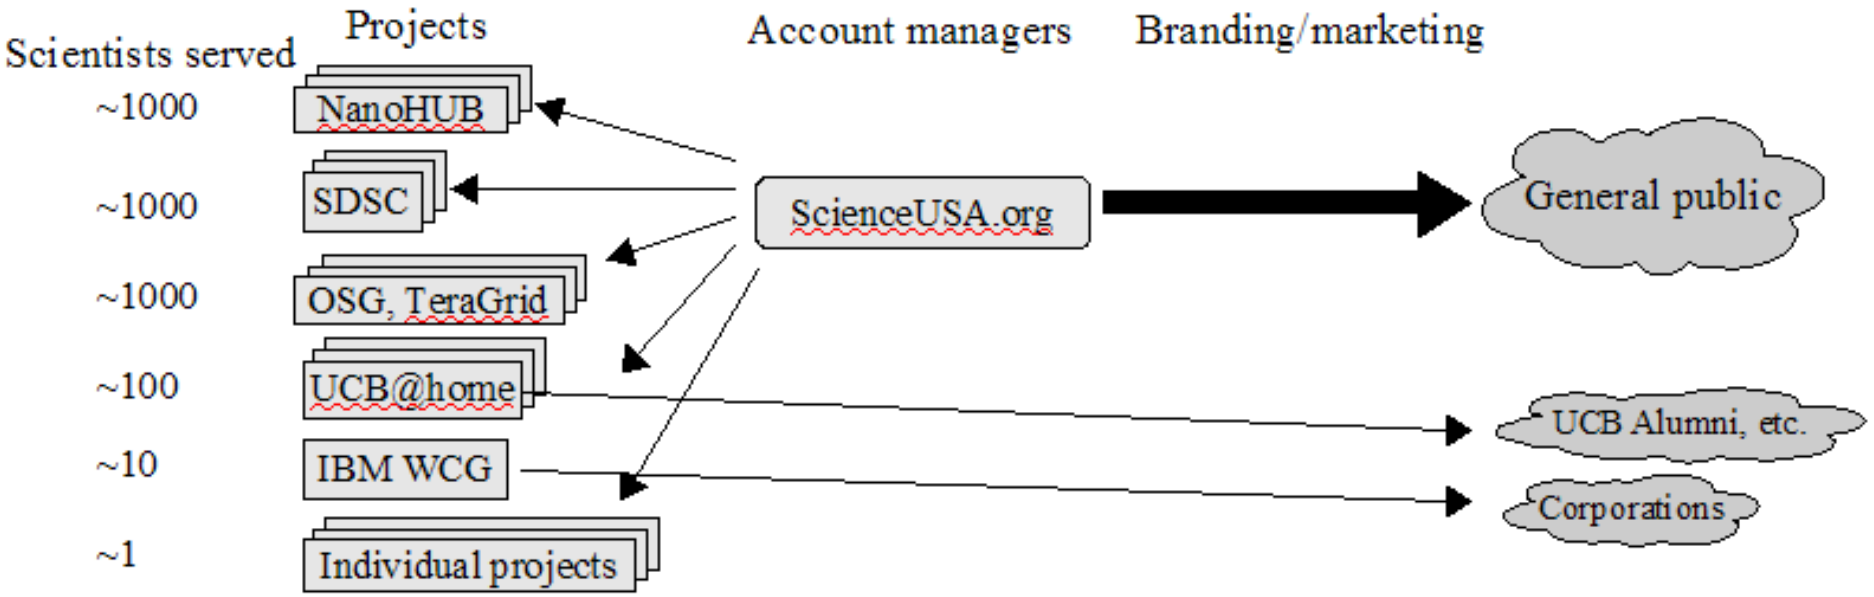


How to realize this?

# Conclusion

- For most scientific computing, volunteer computing is far cheaper than either clouds or clusters
- What's the catch?
  - need to attract volunteers
  - need to learn a new technology (BOINC)
- Related idea: scientific crowd-sourcing
  - Use human brains rather than PCs
  - Stardust@home, GalaxyZoo, etc.